

Analysis of Coppersmith's Block Wiedemann Algorithm for the Parallel Solution of Sparse Linear Systems

ERICH KALTOFEN

Rensselaer Polytechnic Institute
Department of Computer Science
Troy, New York, USA

Outline


- **object representation**
 - black box representations for sparse matrices
- **Wiedemann's algorithm**
 - homogeneous linear systems
- **Coppersmith's blocking**
 - my probabilistic analysis
 - fast solution of singular block-Toeplitz systems
- **our implementation efforts in DSC**
 - timings for test cases (with A. Lobo [DISCO '93])
- **the large sparse linear system challenge**

What is a sparse matrix?

- **matrices with “few” non-zero entries**
 - a band matrix from a finite element method
 - a matrix over GF(2) from integer factoring by the NFS:
52 250 × 50 001 with 1 095 532 entries ≠ 0 (≈ 21/row)
- **matrices with special structure**
 - the Sylvester matrix corresponding to a polynomial resultant

$$R = \begin{bmatrix} a_N & a_{N-1} & \dots & a_0 & & & \\ & a_N & \dots & a_1 & a_0 & & 0 \\ & & \ddots & & \ddots & \ddots & \\ 0 & & & a_N & \dots & \dots & a_0 \\ b_N & b_{N-1} & \dots & b_0 & & & \\ & b_N & \dots & b_1 & b_0 & & 0 \\ & & \ddots & & \ddots & \ddots & \\ 0 & & & b_N & \dots & \dots & b_0 \end{bmatrix}$$

- a “black box” matrix
an efficient program with the specifications



e.g., for the Sylvester matrix R , $R \cdot b$ costs

$$O(N \log N \log \log N)$$

arithmetic operations using fast polynomial multiplication

Symbolic objects given by black box representation are known for many problems:

- symbolic determinants using Gaussian elimination
- the polynomial remainder sequence of $f_0(x)$ and $f_1(x)$ using continued fraction approximations

$$\{q_i(x)\}_{i \geq 2} \quad \text{such that} \quad f_i(x) = f_{i-2}(x) - q_i(x)f_{i-1}(x)$$

- $B^{-1} = P^{-1}U^{-1}L^{-1}$, the LUP factorization of $B \in \mathbb{K}^{N \times N}$.
- streams for infinite objects, such as a program for the i -th order coefficient of a power series

Linear system solution with a black box matrix

Given a black box



$B \in \mathbb{K}^{N \times N}$ singular
 \mathbb{K} an arbitrary field

compute $w \neq \mathbf{0}$ such that $Bw = \mathbf{0}$ “efficiently.”

D. Wiedemann (1986) constructs a Las-Vegas-randomized algorithm that computes w in at most

$3N$ “ $B \cdot b$ steps”

and

$O(N^2)$ additional arithmetic operations in \mathbb{K} .


The algorithm needs $O(N)$ space.

Idea for WIEDEMANN'S algorithm


$B \in \mathbb{K}^{N \times N}$, \mathbb{K} a finite field

$f^B(\lambda) = c'_0 + c'_1 \lambda + \dots + c_{M'} \lambda^{M'} \in \mathbb{K}[\lambda]$ **minimum polynomial** of B :

$$\forall u, v \in \mathbb{K}^N: \forall j \geq 0: u^{\text{tr}} B^j f^B(B)v = 0$$



$$c'_0 \underbrace{u^{\text{tr}} B^j v}_{a_j} + c'_1 \underbrace{u^{\text{tr}} B^{j+1} v}_{a_{j+1}} + \dots + c'_{M'} \underbrace{u^{\text{tr}} B^{j+M'} v}_{a_{j+M'}} = 0$$



$\{a_0, a_1, a_2, \dots\}$ is generated by a linear recursion

Theorem [Wiedemann 1986]: *For random $u, v \in \mathbb{K}^N$,
a linear generator for $\{a_1, a_1, a_2, \dots\}$ is one for $\{I, B, B^2, \dots\}$.*

$$\forall j \geq 0: c_0 a_j + c_1 a_{j+1} + \dots + c_M a_{j+M} = 0$$

\Downarrow (with high probability)

$$c_0 B^j v + c_1 B^{j+1} v + \dots + c_M B^{j+M} v = \mathbf{0}$$

\Downarrow (with high probability)

$$c_0 B^j + c_1 B^{j+1} + \dots + c_M B^{j+M} = \mathbf{0}$$

that is, $f^B(\lambda)$ divides $c_0 + c_1 \lambda + \dots + c_M \lambda^M$

One may compute the c_j from the a_j by the BERLEKAMP/MASSEY algorithm, or by solving a homogenous linear Toeplitz system:

$$\begin{bmatrix} a_M & \cdots & a_1 & a_0 \\ a_{M+1} & a_M & a_2 & a_1 \\ \vdots & & \ddots & \vdots \\ a_{2M-1} & \cdots & & a_{M-1} \end{bmatrix} \underbrace{\begin{bmatrix} c_M \\ c_{M-1} \\ \vdots \\ c_0 \end{bmatrix}}_{\neq 0} = \mathbf{0}$$

for any $M \geq \deg(f^B)$, in particular for $M = N$.

Algorithm *Homogeneous Wiedemann*

Input: $B \in \mathbb{K}^{N \times N}$ singular

Output: $w \neq \mathbf{0}$ such that $Bw = \mathbf{0}$

Step W1: Pick random $u, v \in \mathbb{K}^N$; $b \leftarrow Bv$;
for $i \leftarrow 0$ to $2N - 1$ do $a_i \leftarrow u^{\text{tr}} B^i b$.

$2N$ “ $B \cdot y$ ” steps
 $O(N^2)$ arithm. op’s

Step W2: Compute a linear recurrence generator for $\{a_i\}$,
 $c_L \lambda^L + c_{L+1} \lambda^{L+1} + \dots + c_D \lambda^D$, $L \geq 0, c_L \neq 0$.

$O(N (\log N)^2 \log \log N)$ arithm. op’s

Step W3: $\hat{w} \leftarrow c_L v + c_{L+1} Bv + \dots + c_D B^{D-L} v$;
(With high probability $\hat{w} \neq \mathbf{0}$ and $B^{L+1} \hat{w} = \mathbf{0}$)

Compute first l with $B^l \hat{w} = \mathbf{0}$; **return** $w \leftarrow B^{l-1} \hat{w}$.

$\leq N + 1$ “ $B \cdot y$ ” steps
 $O(N^2)$ arithm. op’s

Other applications of “coordinate recurrences”

- solution of sparse singular inhomogeneous linear systems [K & Saunders 1991]
- processor-efficient parallel poly-log-time algorithms for dense linear systems [K & Pan 1991, 1992]
- Frobenius form computation of an $N \times N$ matrix in $O(N^{2.375})$ field operations [Giesbrecht 1991]
- processor-efficient parallel poly-log-time algorithms for the characteristic polynomial [Eberly 1991, Giesbrecht 1992]
- space-complexity improvement of the Berlekamp polynomial factoring algorithm [K 1991]

Coppersmith's (1992) parallelization (modified)

Use of the block vectors $\mathbf{x} \in \mathbb{K}^{N \times m}$ in place of u
 $\mathbf{z} \in \mathbb{K}^{N \times n}$ in place of v

$$\mathbf{a}_i = \mathbf{x}^{\text{tr}} B^{i+1} \mathbf{z} \in \mathbb{K}^{m \times n}$$

Find a vector polynomial $c_L \lambda^L + c_{L+1} \lambda^{L+1} + \dots + c_D \lambda^D \in \mathbb{K}^n[\lambda]$,
such that

$$\forall j \geq 0: \sum_{i=L}^D \mathbf{a}_{j+i} c_i = \sum_{i=L}^D \mathbf{x}^{\text{tr}} B^{i+j} B \mathbf{z} c_i = \mathbf{0} \in \mathbb{K}^{m \times n}$$

Then, analogously to before, with high probability

$$\widehat{\mathbf{w}} = \sum_{i=L}^D B^{i-L} \mathbf{z} c_i \neq \mathbf{0}, \quad B^{L+1} \widehat{\mathbf{w}} = \sum_{i=L}^D B^i B \mathbf{z} c_i = \mathbf{0} \in \mathbb{K}^N$$

Questions

- how many \mathbf{a}_i are needed?
- how can the c_i be computed from the \mathbf{a}_i ?
- probability of success? does method work on “pathological” B 's?

- iteration length: $i \leq \frac{N}{m} + \frac{N}{n} + \frac{2n}{m} + 1$
- computation of c_i :
 - Coppersmith has generalized the Berlekamp/Massey algorithm
 - Block-Toeplitz solver by Kailath et al. (1979)
Also yields the rank of B .

Both methods require $O((m+n)N^2)$ arithmetic operations in \mathbb{K} or with $m+n$ processors $O(N^2)$ parallel time.

- Divide-and-conquer Toeplitz-like solver [Bitmead & Anderson 1980, Morf 1980, K. 1993]:

Requires $O((m+n)^2 N (\log N)^2 \log \log N)$ arithmetic ops. in \mathbb{K} .

Probabilistic analysis

Theorem: *If B is singular with*

$$\deg f^B = 1 + \text{rank } B \quad (1)$$

then we find $w \neq \mathbf{0}$ with $Bw = \mathbf{0}$ for random \mathbf{x}, \mathbf{z} with probability

$$\geq 1 - \frac{1 + 2 \text{rank } B}{\text{card } \mathbf{K}} \geq 1 - \frac{2N - 1}{\text{card } \mathbf{K}}$$

Condition (1) can be enforced by “randomly mixing” B à la Beneš/Wiedemann (1986), for instance

$\tilde{B} \leftarrow V \cdot B \cdot W \cdot G$ where V realizes random row permut. network
 W realizes random col. permut. network
 G is random diagonal

Parallel coarse-grain realization

The ν^{th} processor computes the ν^{th} column of \mathbf{a}_i , $i \approx \frac{N}{m} + \frac{N}{n}$

Running-time comparisons

	Wiedemann	seq. blocked W.	par. blocked W.
		m, n fixed $\varepsilon = \frac{n}{m} + \frac{1}{n} < 1$	n processors $\frac{m}{n}$ fixed
# of $B \cdot y$ products	$2N$	$3N$	$(1 + \varepsilon)N + 2$
# of arithm. operations	$O(N^2)$	$O(N^2)$	$O\left(\left(1 + \frac{\log N}{n}\right)N^2\right)$ (parallel w/o FFT) $O(n^2 N (\log N)^2)$ $\times \log \log N$ (sequent. w. FFT)
amount of storage	$O(N^2)$	$O(N)$	$O_\varepsilon(N)$

Proof idea for probabilistic analysis

$\sum_{i=0}^{N/n} B^{i+1} z c_i = \mathbf{0}$ is a linear condition on $c_i \equiv$ block-Krylov system

If $\text{rank}(\text{block-Toeplitz}) = \text{rank}(\text{block-Krylov})$
then every solution of block-Toeplitz system is one for block-Krylov system

The *generic* block-Toeplitz/Krylov systems can be specialized to the *generic* Toeplitz system of the Wiedemann algorithm

Thus the rank condition holds for the generic systems, hence for the ones obtained at random specializations by the Schwartz/Zippel lemma.

DSC features

- uses low level UNIX IP/TCP/UDP process communication
- can distribute C and Lisp source code
- network can be interactively monitored for progress and faults
- computers are auto-selected w.r.t. local resources and work load implemented by A. Diaz and M. Hitz [DISCO '93]
- supports co-routine calling mechanism implemented by A. Diaz [DISCO '93]
- messages are digitally signed for secure communication

Test 1: sparse random matrices over GF(32 749)

N	Task	Blocking Factor		
		2	4	8
10 000 [†]	(1) $\langle \mathbf{a}_i \rangle$	7:29	3:54	2:09
	(2) b-massey	2:25	4:08	8:00
	(3) evaluation	3:47	1:59	1:05
	total	13:41	10:06	11:14
20 000 [‡]	(1) $\langle \mathbf{a}_i \rangle$	57:17	28:43	15:21
	(2) b-massey	9:48	16:36	33:39
	(3) evaluation	29:42	14:44	7:53
	total	96:47	60:02	56:53

CPU Time for different blocking factors in *hours:minutes*
each processor rated at 28.5 MIPS

† \approx 350 000 non-zero entries

‡ \approx 1 300 000 non-zero entries

Test 2: sparse matrices over GF(2)

N	Task	Blocking Factor		
		1×32	2×32	3×32
$20\,000^\ddagger$	(1) $\langle \mathbf{a}_i \rangle$	1:12	0:40	0:30
	(2) b-massey	0:25	0:31	0:39
	(3) evaluation	0:29	0:28 $^\diamond$	0:10
	total	2:06	1:39	1:19
$52\,250^*$	(1) $\langle \mathbf{a}_i \rangle$	3:53	2:11	1:37
	(2) b-massey	2:30	3:09	3:54
	(3) evaluation	1:15	0:33	0:22
	total	7:38	5:53	5:53

CPU Time for different blocking factors in *hours:minutes*

32 bit operations performed simultan. as one computer word op.

$\ddagger \approx 1\,300\,000$ non-zero entries picked at random

* from NFS integer factoring; $\approx 1\,100\,000$ non-zero entries

note: unblocked Wiedemann algorithm takes ≈ 111 hours

\diamond first $\hat{w} = \mathbf{0}$.

Test 3: very large sparse matrix over GF(2)

N	Task	Blocking Factor		
		1×32	2×32	3×32
100,000*	(1) $\langle a^{(i)} \rangle$	77:37	44:05	27:28
	(2) b-massey	10:03	12:28	15:42
	(3) evaluation	74:37	27:48	11:09
	total	162:17	84:31	54:19

CPU Time for different blocking factors in *hours:minutes*
each processor rated at 28.5 MIPS

32 bit operations performed simultan. as one computer word op.

★ 10 304 243 non-zero entries picked at random

The large sparse linear system challenge

Solve a sparse $100\,000 \times 100\,000$ linear system over $\text{GF}(2^{32} - 5)$ with 10 000 000 non-zero entries.

Note Amdahl's law:

$$\frac{T_{\text{par}}}{T_{\text{seq}}} = \alpha + \frac{1 - \alpha}{p(1 - c)} \quad \text{where } p \text{ \# of processors}$$

α ratio spent in seq. part

c ratio of par. part spent in commun.